Knowledge Reuse: Temporal-Abstraction Mechanisms for the Assessment of Children's Growth

Manon M. Kuilboer, MD*, Yuval Shahar, MD*,
Darrell M. Wilson, MD** Mark A. Musen, MD, PhD *

*Section on Medical Informatics and ** Department of Pediatrics
Stanford University School of Medicine
Stanford, CA 94305-5479

ABSTRACT

Currently, many workers in the field of medical informatics realize the importance of knowledge reuse. The PROTÉGÉ-II project seeks to develop and implement a domain-independent framework that allows system builders to create custom-tailored role-limiting methods from generic reusable components. These new role-limiting methods are used to create domain- and task-specific knowledge-acquisition tools with which an application expert can generate domain- and task-specific decision-support systems. One required set of reusable components embodies the problem-solving knowledge to generate temporal abstractions. Previously, members of the PROTÉGÉ-II project have used these temporal-abstraction mechanisms to infer the presence of myelotoxicity in patients with AIDS. In this paper, we show that these mechanisms are reusable in the domain of assessment of children's growth.

1. REUSE OF KNOWLEDGE

The medical-informatics community has acknowledged the importance of the development of approaches that will make it possible to reuse components of knowledge. These components may be reusable lexicons, ontologies, tasks, or problem-solving methods [1]. Knowledge-acquisition (KA) tools that allow developers to reuse predefined models of problem-solving methods impose method-specific architectures on the knowledge bases that the tools are used to construct [2]. These architectures are often well suited for the performance of only a narrow range of tasks. It takes considerable time to develop new problem-solving methods, as well as the associated knowledge-acquisition tools and run-time environments. Consequently, there is little chance that researchers can reuse one another's work when the abstract definitions of the method, the KA tools, and the run-time system must be integrated completely. One type of potentially separable, reusable component is the problem-solving method itself, which is an abstract strategy to apply in new domains. In PROTÉGÉ-II, problem-solving methods either can or cannot be decomposed into subtasks. Those problem-solving methods that are not decomposable are called *mechanisms*. A problem-solving mechanism solves one particular problem for a task. It can be used for the same task in different domains, because it allows separation of domain-dependent knowledge (e.g., a domain-specific ontology) from domain-independent problem-solving knowledge. The latter can be task-specific.

Temporal-abstraction mechanisms are an example of problem-solving methods that are domain independent [2]. These mechanisms have been applied to the domain of inferring myelotoxicity by the software system RÉSUMÉ [4]. In this paper, we show that the temporal-abstraction mechanisms are indeed reusable by describing how RÉSUMÉ can perform the task of temporal abstraction for the assessment of children's growth.

2. METHODS AND ASSUMPTIONS

We shall introduce the domain, and give a brief overview of the temporal-abstraction mechanisms as well as our knowledge-acquisition cycle.

2.1 Assessment of children's growth.

The task of assessing of the normality of children's growth is suitable for use of temporal-abstraction mechanisms, in that the task requires analysis of time-dependent data. The main difference from the domain of myelotoxicity is that the time scale is much larger (years, instead of days). A second difference is that only a few data points are available (3 to 15 data points; usually about 5).

In practice, a child's height, maturation, and other relevant clinical data are recorded in the patient chart at each visit. The height and maturation data also are noted on a growth chart. A *growth chart* is a printed set of standard growth curves that represents the mean and the 95-percent confidence interval of the expected height of a child at a particular age. The physician then assesses the normality of the growth of a child, taking into consideration the height of the parents, as well as the maturation of the child.

2.2 Temporal-Abstraction Mechanisms

Our model of temporal abstraction is a point-based, discrete, closed-interval temporal model [3]. Time intervals have associated parameter values. Parameters can be either *primitive* (e.g., raw data, such as "HEIGHT = 170 cm"), or *abstract* (e.g., characterizations such as "rate of change in the HTSDS = FAST," where "HTSDS" denotes the height expressed as standard deviation score). If an interval has any abstraction attached to it, we refer to it as an *abstracted interval.* We distinguish among three types of abstractions: *state*, *gradient*, and *rate* abstractions. These types correspond, respectively, to the classification of the magnitude of the parameter, to the sign of the derivative of the parameter during the interval, and to the absolute magnitude of the derivative of the parameter during the interval (e.g., LOW, DECREASING, and FAST abstractions, respectively, for the HTSDS parameter). We shall review each of the three temporal abstraction mechanisms that we have applied to the task of assessing children's growth.

Temporal Point Abstraction. This mechanism abstracts into a single class one or more primitive or abstract parameter values. The output of this mechanism is an abstracted state, such as a classification "maturation = AVERAGE", that might be the classification of the combined values of several Tanner stages expressed as standard deviation score (SDS) (maturation parameters), and a bone-age measurement expressed as SDS of one point in time.

Temporal Inference. This mechanism infers domain-dependent sound logical conclusions over a single interval or two meeting intervals, such as when the task allows us to concatenate two meeting intervals with a LOW value of the HTSDS state abstraction into one (HTSDS = LOW) interval. Such abstractions have the *concatenable* semantic property [3]. If abstractions are concatenable, temporal inference also determines the *value* of the joined abstraction (e.g., DECREASING and INCREASING might be concatenated into NONMONOTONIC).

Temporal Interpolation. This mechanism bridges nonmeeting temporal points or intervals. Examples of its use include creating a (HTSDS = LOW) episode from two close abstracted points, or bridging a gap between two known intervals of a (HTSDS = DECREASING) gradient abstraction. The temporal-interpolation mechanism uses a domain-specific function, the *Δ-function*, that returns the maximal allowed time gap between the two intervals that still enables interpolation over the gap, joining the two intervals and the gap into one interval with an abstraction defined by the appropriate abstraction-inference table. For instance, in any context, joining two intervals where the state abstraction of the HTSDS parameter was classified as LOW into a longer interval where the state abstraction of the HTSDS parameter is classified as LOW, might depend on the time gap separating the two intervals, on the properties of the HTSDS state abstraction for the value LOW in that context, and on the lengths of time during which the LOW property was known both before and after the time gap.

The temporal-abstraction mechanisms do not operate in a fixed order, but instead iterate alternately, activated by the currently available data and the previously derived abstractions. A truth-maintenance system updates the temporal-interval conclusions and propagates changes caused by any updates, since these conclusions are by nature *defeasible*, their validity being dependent on primitive data that might be modified [4].

The growth-chart domain-specific knowledge required by the temporal-abstraction mechanisms (e.g., classification tables for point abstraction) is represented as a *growth-chart parameter-properties ontology*—that is, as a pictorial theory that represents the raw and abstract data in this domain (e.g., HTSDS state abstractions, HTSDS rate abstractions, bone-age maturation stages), and the relations among them. The structure of the growth-chart domain parameter ontology, of which an example will be shown in Figure 1, is a combination of two graphs. The first is an IS-A graph that represents the relations of abstractions, subclasses, and subcontexts (e.g., HTSDS-CLASSIF is a state abstraction of the parameter HTSDS). The second is an INTO hierarchy that represents which one or more parameters map into which other parameter (e.g. HEIGHT, GENDER, AGE, and POPULATION DATA map into HTSDS).

2.3 The Knowledge-Acquisition Cycle

Most of our KA work has concentrated on creating a proper structure for the growth-chart domain parameter-properties ontology, and instantiating the class slots of the ontology (e.g., by instantiating specific temporal-abstraction functions for particular parameter classes) with the help of a domain expert. In forming abstractions, we focused on the problem of distinguishing between normal and abnormal growth. The expert for the growth-chart domain was one of us (DW), a pediatric endocrinologist.

First, we conducted a meeting in which we explained the principles of the temporal-abstraction mechanisms. We provided DW with literature about this methodology, so that he would understand how we wished to abstract his knowledge of the growth-chart domain.

The following KA cycle was efficient, because the temporal-abstraction mechanisms clearly define the different kinds of knowledge that we had to acquire for the temporal-abstraction task. The mechanisms thus define what McDermott [5] refers to as *knowledge roles*. First, we examined the mechanisms to itemize the knowledge roles (e.g., to identify a trend in the SDS of height, the ontology would need the parameter height and knowledge about the population distribution of height). Second, we constructed a structure for the ontology that satisfied those roles (e.g., we had to define that HEIGHT has a state abstraction, but neither a gradient abstraction nor a rate abstraction). Third, we acquired the attributes and functions of the parameters in the ontology (e.g., what is the minimal difference between two measurements before they are considered to be different, or which trends need to be recognized and how they change by age).

3. RESULTS

We shall describe the result of our parameter-properties ontology, as well as our experience with the application of RÉSUMÉ to the growth-chart domain.

3.1 The Parameter-Properties Ontology for the Growth-Chart Domain

The main structure of the ontology — the division between abstract and primitive parameters — is shown in Figure 1. The abstractions are subdivided into state, gradient, and rate abstractions. All parameters are either a primitive parameter or one of these abstractions, as shown in the structure. We discuss in Section 3.2 the asymmetric shape of the ontology around the parameter "HEIGHT expressed as SDS" (HeightSD abstraction). We changed the structure of the ontology many times. An important reason for these changes was that a particular variable would or would not be found to be necessary for the task of assessing normality of growth, as dictated by the temporal-abstraction mechanisms and the expert. For example, we removed data on normality of X-ray examinations of the skeleton, except for those from the X-ray examination of the left hand for boneage; laboratory values for indicators of illnesses in internal medicine; and weight. As we learned from DW, these data are important in diagnosing a disorder, but are not relevant in distinguishing normal versus abnormal growth of a child from abstracted intervals as created by temporal abstraction.

3.2 The Application of RÉSUMÉ to the Growth-Chart Domain.

With the acquired knowledge, we were able to test RÉSUMÉ for the growth-chart domain on three case examples — one patient with precocious puberty, and two patients who showed a genetically short height, but otherwise normal growth. All three cases yielded intervals from which MK could draw as the conclusion the diagnosis in the patient record. An example of the results of these preliminary tests is shown in Figure 2.

Many more intervals were created by RÉSUMÉ than those shown in Figure 2, but the user does not need to examine all of them. RÉSUMÉ uses the intervals to

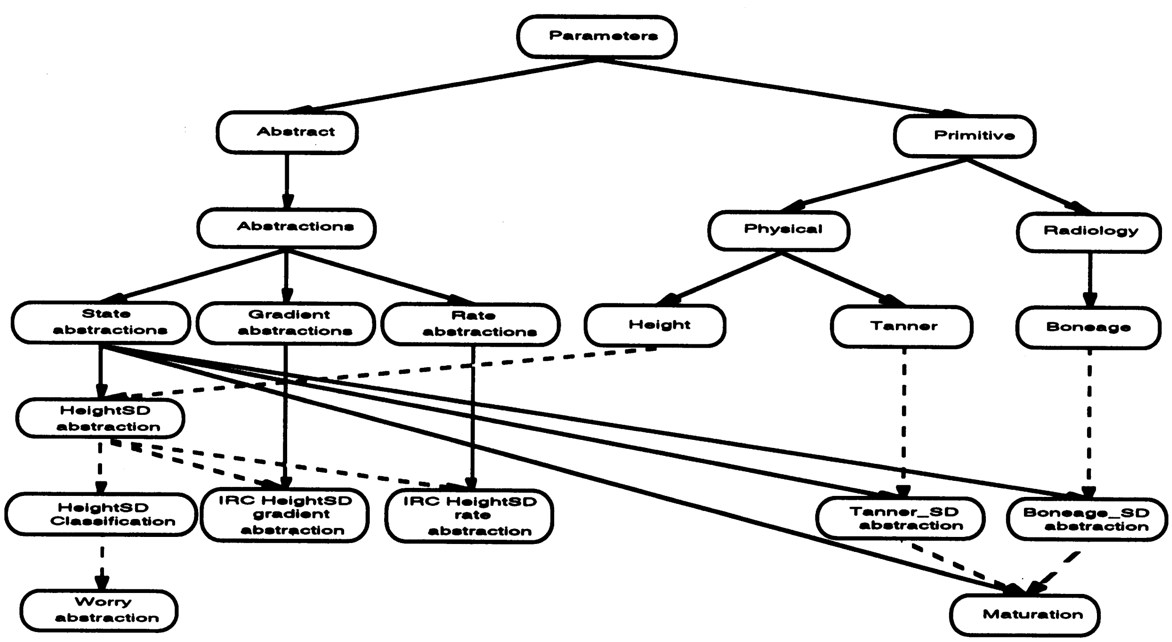

Figure 1: *Part of the structure of the growth-chart parameter ontology. A combination of two graphs is shown: IS-A graph = (—>), showing the IS-A relations between objects, and an INTO hierarchy = (-->) that shows which parameters map into which other parameter.*

create the larger abstracted intervals. By treating the resultant complete collection of abstractions as a database, we could sift through the created intervals using RÉSUMÉ's temporal query module to select only those intervals that answer specific questions. In our case, we wanted to examine those intervals that match the criteria for normal or abnormal growth as specified by DW. These criteria are not fixed and can be set or fine-tuned by the individual user. In all three of our test cases, the result of our queries yielded the same conclusion about normality of the growth as that reached by DW. We thus conclude that RÉSUMÉ generates those intervals necessary to reach the correct interpretations for the three test cases.

4. DISCUSSION

The KA cycle gave us several important insights. The first is that the system builders and DW had different mental models of the knowledge necessary for the growth-assessment task. We learned that there is a delicate balance between adhering too closely to the formal temporal-abstraction mechanisms model and adapting too much to the clinician's intuition and vocabulary. If the approach is too formal or too clinical, either the expert refuses to continue to collaborate, or the knowledge engineer does not get the information she needs to instantiate the method with which she chose to work.

The harmonization of our mental models [6] progressed slowly. We learned from each other, and were finally able to speak each other's language well enough to instantiate the slots of the parameter-properties ontology. Asking an expert to fill out the tables manually without any support would be unrealistic. Also, simply trying to enforce the formal model might annoy the expert or even cause him to lose confidence in the methodology, especially when the knowledge engineer is working on the edge of what the expert is willing to quantify explicitly. For example, cut-offs for classifications were generally not difficult for DW to give. However, DW did find it difficult to specify those cutoffs that, in his model, were closely related to other variables that together could influence the interpretation.

One conclusion is that an automated KA interface that enables the expert to modify the knowledge base in his own time, and that provides some support for the KA, might enhance the efficiency of the KA cycle, even though it cannot completely replace a human-supported cycle. We are currently developing systems that create KA tools automatically from domain models in the PROTÉGÉ-II project. This project aims to provide tools with which developers can generate KA tools for specific domains with a specific problem-solving method. These KA tools can in turn generate decision-support systems, specialized for a specific task.

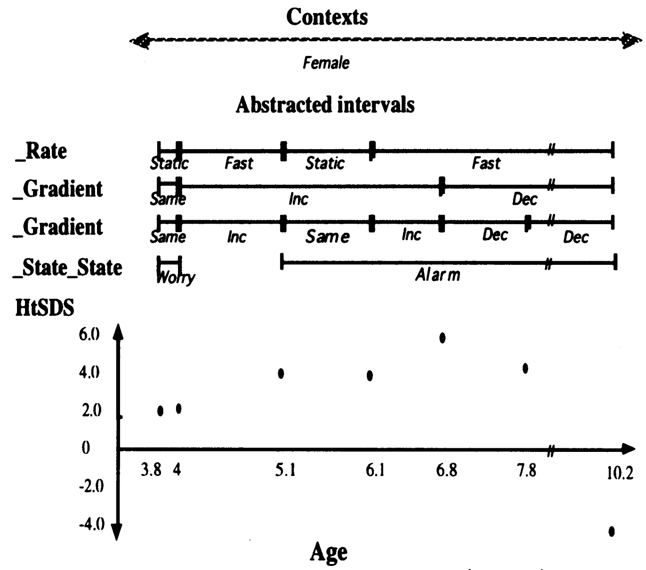

Figure 2: *Part of the result of running RÉSUMÉ using the acquired parameter ontology on a test case (female patient with precocious puberty). Time points, intervals, and associated parameter values are shown, starting with the HTSDS parameter. Next, the state abstractions of the state abstractions of HTSDS are shown, with values WORRY and ALARM. The higher two levels show two stages of the gradient abstractions with values SAME, INC= INCREASING, DEC=DECREASING. The rate abstraction shows the values FAST and STATIC.*

Our preliminary testing showed that the temporal-abstraction mechanisms as used in RÉSUMÉ yield the abstractions needed to determine whether the growth curve of a patient is normal. We draw our conclusions with care because we have been able to test only three cases in the growth domain. The results suggest that the RÉSUMÉ-created abstracted intervals resemble the interval-based abstractions that physicians use, which these physicians normally do not make explicit. In addition, it is difficult for physicians to combine all available abstractions quickly (e.g., HTSDS, TannerSDS, and BoneageSD), whereas RÉSUMÉ can do so easily.

The domain of children's growth is different from the protocol-treatment domains on which RÉSUMÉ has been tested previously [7, 8]. The scale of the timeline in the growth-chart domain is large compared with that of typical patient-visit oriented domains. This large scale places different constraints on the required tables and gap functions. For some of the maximal-gap functions, we were able to use a simple truth-persistence assumption (i.e., an infinite maximal-gap function, which means that, given some value at some point in time, we assume that it stays the same unless there is a reason to change that assumption). This assumption has not been appropriate for the other domains in which RÉSUMÉ has been tested, because the primary data values could fluctuate much more widely. Growth-chart interpretations require mainly high-level abstractions (e.g., the classification of HTSDS is LOW), rather than low-level abstractions (e.g., HTSDS is equal to 0.6) that tend to fluctuate wildly, which, just as in many other domains, makes the simple truth-persistence assumption unrealistic. Using high-level abstractions (e.g., the classification of the HTSDS) reduced variations to the point of enabling us to use an infinite gap function as a default.

In the growth-chart domain, the absolute-height value itself plays only a minor role in the interpretation of the growth curve. The *abstractions* of height and maturation measures are more important. For example, we use HTSDS as only a state abstraction of HEIGHT, but we do not create gradient or rate abstractions of HEIGHT. This example can be seen in the parameter ontology of the growth-chart domain, where height is abstracted into HTSDS, but not into HEIGHT_GRADIENT or HEIGHT_RATE. On the other hand, for the HTSDS parameter, all three abstraction types are of interest; consequently, they are represented in the ontology (HTSDS_STATE, HTSDS_GRADIENT, and HTSDS_RATE). This example is one where the process of creating the ontology made knowledge explicit that previously was only implicit.

It is remarkable that, in this seemingly simple domain, we encountered symbolic state abstractions of symbolic state abstractions (e.g., the state abstraction ALARM of the state abstraction VERYHIGH of HTSDS), as well as computational numeric and symbolic transformations that are four levels deep (e.g., from the Tanner stage to the maturation score). The creation of these abstractions is in contrast with most domains that we have studied, in which symbolic abstractions are created only from numeric variables. Although some of these transformations lose information, they focus the interpretation process. RÉSUMÉ maintains (and answers queries about) the original, as well as intermediate, data and abstractions.

In our experience, the temporal-abstraction mechanisms provide efficient guidelines for what type of knowledge (in which roles) we must acquire in the interaction with an expert, to perform a certain task (e.g., temporal abstraction). This guidance during one KA process is a strong advantage of reusable mechanisms, which are themselves domain independent.

Because the parameter ontology that results from one KA project exists separately from the temporal-abstraction mechanisms, the ontology is itself reusable. For example, in the work of Kohane and Haimowitz on automated trend deduction with alternate temporal hypotheses [9], the two components are not separated, which makes it difficult to reuse incorporated knowledge. Our parameter ontology, on the other hand, could be reused for tasks such as diagnosing a disorder, making a treatment suggestion, or predicting a growth trajectory.

Acknowledgments
Dr.Kuilboer was supported by the following Dutch organizations: Ter Meulen Fund of the Royal Netherlands Academy of Arts and Sciences, Genootschap Noorthey, the Foundation 'De Drie Lichten' and the VSB-Foundation. Dr. Musen is recipient of NSF Young Investigator Award IRI-9257578. We also acknowledge the support of the Fulbright Commission, the Agency for Health Care Policy and Research (under grant HS06330) and the National Library of Medicine (under grants LM05157 and LM05305). Lyn Dupre provided helpful comments on a previous draft of this paper.

References

1. Musen MA. Dimensions of knowledge sharing and reuse. *Computers and Biomedical Research.* **25**:435–467, 1992.
2. Musen MA, Tu SW. Problem-solving models for generation of task-specific knowledge-acquisition tools. *Knowledge-Oriented Software Design.* (J. Cuena, ed.), Amsterdam: Elsevier, 23-49, 1993.
3. Shoham Y. Temporal logics in AI: Semantical and ontological considerations. *Artificial Intelligence.* **33**:89–104, 1987.
4. Shahar Y, Tu SW, Musen MA. Knowledge acquisition for temporal-abstraction mechanisms. *Knowledge Acquisition.* **4**:217–236, 1992.
5. McDermott J. Preliminary steps toward a taxonomy of problem-solving methods. In: *Automating Knowledge Acquisition for Expert Systems.* (Marcus S, ed.), Boston, MA: Kluwer Academic. pp. 225–255, 1988.
6. Regoczei S, Plantinga EPO. Creating the domain of discourse: Ontology and inventory. *International Journal of Man–Machine Studies.* **27**:235–250, 1987.
7. Shahar Y, Musen MA. *RÉsumÉ:A temporal-abstraction system for patient monitoring.* Computers and Biomedical Research 1993;26(3):255-273.
8. Shahar Y, Tu SW, Musen MA. Using temporal-abstraction mechanisms for patient monitoring. *AAAI Symposium on Artificial Intelligence in Medicine.* Stanford University. pp. 95–99, 1992.
9. Haimowitz IA., Kohane IS. Automated trend deduction with alternate temporal hypotheses. IJCAI, 1993.